



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/808,561	03/24/2004	Ofir Zohar	ASSIA 21.061	4776
26304 7590 12/12/2007 KATTEN MUCHIN ROSENMAN LLP 575 MADISON AVENUE NEW YORK, NY 10022-2585			EXAMINER BRADLEY, MATTHEW A	
			ART UNIT 2187	PAPER NUMBER
			MAIL DATE 12/12/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

mn

Office Action Summary	Application No.	Applicant(s)	
	10/808,561	ZOHAR ET AL.	
	Examiner	Art Unit	
	Matthew Bradley	2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 01 October 2007.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-33 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-33 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☒ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Response to Amendment

This Office Action has been issued in response to amendment filed 1 October 2007. Applicant's arguments have been carefully and fully considered but they are not persuasive. Accordingly, this action has been made FINAL.

Claim Status

Claims 1-33 remain pending and are ready for examination.

Specification

The specification is objected to as failing to provide proper antecedent basis for the claimed subject matter. See 37 CFR 1.75(d)(1) and MPEP § 608.01(o). Correction of the following is required: Newly amended independent claims **1, 14, 26, 28, 30, and 32** recite in part, caches -"configured to" or "adapted to"- "function" or "operate" – as controllers. Insofar as it appears to be clear, the Examiner is having difficulty locating explicit support in the specification for the caches operating as controllers as explained in further detail *infra*.

Claim Rejections - 35 USC § 112

The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

Claims **1-33** are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to

Art Unit: 2187

one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention. Newly amended independent claims **1, 14, 26, 28, 30, and 32** recite in part, caches -"configured to" or "adapted to"- "function" or "operate" – as controllers. Insofar as it appears to be clear, the Examiner is having difficulty locating explicit support in the specification for the caches operating as controllers. Referring back to Applicant's specification, page 16 line 23 to page 17 line 11, the following is found, 'caches 20 are coupled to interfaces 26 by any suitable fast coupling system known in the art, such as a bus or a switch, so that each interface is able to communicate with, and transfer data to and from, any cache. Herein the coupling between caches 20 and interfaces 26 is assumed, by way of example, to be by a first cross-point switch 14. Interfaces 26 operate substantially independently of each other. ***Caches 20 and interfaces 26 operate as a data transfer system 27, transferring data between hosts 52 and disks 12.***' At least as is shown in the specification, the caches themselves are not operating or functioning as controllers. It appears as though the interfaces, as recited in the specification, work with the caches to achieve the functionality of a controller, not the caches just of themselves. Accordingly, claims 1, 14, 26, 28, 30, and 32 fail to comply with the written description requirement.

Any claim not specifically addressed is rejected to at least by virtue of its dependency.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-8, 12-21 and 25 are rejected under 35 U.S.C. 102(e) as being anticipated by Hicken et al. (US 2004/0153727).

With respect to claim 1, Hicken et al. disclose a method for managing a data storage system (**300 of Fig. 3; paragraph 0038, lines 4-10**), comprising:

- configuring a first cache (**339 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at a first range of logical addresses (LAs) in a storage device (**paragraph 0038, lines 13-17; paragraph 0039, lines 7-10; paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2**);
- configuring a second cache (**333 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (**paragraph 0039, lines 5-7; paragraph 0041, lines 10-14**);
- configuring one or more third caches (**338 of Fig. 3**) to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device the first, second, and third caches being implemented in three separate physical units and adapted to function as controllers substantially independently of each other; (**paragraph 0041, lines 14-17**);

- detecting an inability of the second cache to retrieve data from or store data at the first range of LAs (**paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails as well**); and
- reconfiguring at least one the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (**paragraph 0044**).

With respect to claim 2, Hicken et al. disclose the method according to claim 1, and comprising configuring one or more interfaces (**CPUs 331 and 336 of Fig. 3**) to receive input/output (IO) requests (**paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16**) from host processors (**310 of Fig. 3; paragraph 0038, lines 4-7**) directed to specified LAs (**paragraph 0025, lines 3-6**) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (**paragraph 0039, lines 15-16**).

With respect to claim 3, Hicken et al. disclose the method according to claim 2, wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (**paragraph 0038, lines 13-17**), and wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response

to the mapping (**paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches**), and wherein detecting the inability comprises generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (**paragraph 0042, lines 9-13**), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (**paragraph 0042, lines 9-13**).

With respect to claim 4, Hicken et al. disclose the method according to claim 1, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises processing data in the first cache and the one or more third caches so as to restore the first cache and the one or more third caches to a state of full data redundancy (**paragraph 0044, lines 15-24**).

With respect to claim 5, Hicken et al. disclose the method according to claim 4, wherein processing the data comprises classifying data in the first cache into a plurality of data groups (**paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified**).

With respect to claim 6, Hicken et al. disclose the method according to claim 5, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the one or more third caches (**paragraph 0043, lines 10-15; dirty data is stored on the third cache 338**).

With respect to claim 7, Hicken et al. disclose the method according to claim 5, wherein one of the data groups comprises dirty data, and wherein processing the data

comprises storing the dirty data at the storage device (**paragraph 0043, lines 10-15; dirty data is flushed to the storage units**).

With respect to claim 8, Hicken et al. disclose the method according to claim 1, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (**paragraph 0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1**).

With respect to claim 12, Hicken et al. disclose the method according to claim 1, and comprising providing a system manager (**host computer 310 and CPUs 331 and 336 of Fig. 3**) which is adapted to configure the first, second and one or more third caches (**paragraph 0025, lines 4-6; paragraph 0039, lines 15-16**), to detect the inability (**paragraph 0028, lines 1-4**), and to reconfigure the at least one of the first cache and the one or more third caches (**paragraph 0042, lines 9-13, lines 18-22**).

With respect to claim 13, Hicken et al. disclose the method according to claim 12, wherein providing the system manager comprises incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (**host computer 310 and CPUs 331 and 336 of Fig. 3**), and operating the one or more manager processing units in a cooperative manner (**paragraph 0040; all of the CPUs are connected and work together**).

With respect to claim 14, Hicken et al. disclose a data storage system, comprising:

- a storage device (**300 of Fig. 3; paragraph 0038, lines 4-10**) wherein data is stored at logical addresses (LAs);
- a first cache (**339 of Fig. 3**) which is configured to perform at least one of the operations of retrieving data from and storing data at a first range of LAs in the storage device (**paragraph 0038, lines 13-17; paragraph 0039, lines 7-10; paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2**);
- a second cache (**333 of Fig. 3**) which is configured to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (**paragraph 0039, lines 5-7; paragraph 0041, lines 10-14**);
- one or more third caches (**338 of Fig. 3**) which are configured to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device the first, second, and third caches being implemented in three separate physical units and adapted to function as controllers substantially independently of each other; (**paragraph 0041, lines 14-17**); and
- a system manager (**host computer 310 and CPUs 331 and 336 of Fig. 3**)
 - which is adapted to detect an inability of the second cache to retrieve data from or store data at the first range of LAs (**paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails as well**), and

- which reconfigures at least one the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (**paragraph 0044**).

With respect to claim 15, Hicken et al. disclose the storage system according to claim 14, and comprising one or more interfaces (**CPUs 331 and 336 of Fig. 3**) which are configured to receive input/output (IO) requests (**paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16**) from host processors (**310 of Fig. 3; paragraph 0038, lines 4-7**) directed to specified LAs (**paragraph 0025, lines 3-6**) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (**paragraph 0039, lines 15-16**).

With respect to claim 16, Hicken et al. disclose the storage system according to claim 15, wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (**paragraph 0038, lines 13-17**), and wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping (**paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches**), and wherein detecting the inability comprises the system

manager generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (**paragraph 0042, lines 9-13**), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (**paragraph 0042, lines 9-13**).

With respect to claim 17, Hicken et al. disclose the storage system according to claim 14, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the first cache processing data therein and the one or more third caches processing data therein so as to restore the first cache and the one or more third caches to a state of full data redundancy (**paragraph 0044, lines 15-24**).

With respect to claim 18, Hicken et al. disclose the storage system according to claim 17, wherein processing the data comprises classifying data in the first cache into a plurality of data groups (**paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified**).

With respect to claim 19, Hicken et al. disclose the storage system according to claim 18, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the one or more third caches (**paragraph 0043, lines 10-15; dirty data is stored on the third cache 338**).

With respect to claim 20, Hicken et al. disclose the storage system according to claim 5, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the storage device (**paragraph 0043, lines 10-15; dirty data is flushed to the storage units**).

With respect to claim 21, Hicken et al. disclose the storage system according to claim 1, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (**paragraph 0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1).**

With respect to claim 25, Hicken et al. disclose the storage system according to claim 14, wherein the system manager comprises one or more manager processing units incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (**host computer 310 and CPUs 331 and 336 of Fig. 3**), and wherein the one or more manager processing units operate in a cooperative manner (**paragraph 0040; all of the CPUs are connected and work together**).

Claims 26-33 are rejected under 35 U.S.C. 102(e) as being anticipated by Henry et al. (US 6,898,666).

With respect to claim 26, Henry et al. teach a storage system, comprising:

- one or more mass storage devices, coupled to store data at respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23, lines 55-67**);
- a plurality of interim fast-access-time caches (**cache pools 1 and 2**), configured to operate as controllers substantially independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (**column 5, lines 45-55**); and coupled to receive data from and provide data to

the one or more -mass-storage devices having LAs within the respective second range (**column 4, lines 55-59**); and

- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).
- a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails (**paragraph 0044**).

With respect to claim 27, Henry et al. teach the storage system according to claim 27, wherein the mass storage devices comprise one or more disks (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

With respect to claim 28, Henry et al. teach a method for storing data, comprising:

- storing the data in one or mass storage devices having respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23 & 55-67**);
- assigning to each of a plurality of interim-fast-access-time nodes (**cache pools 1 and 2**), configured to operate as controllers substantially independently of one another, a respective second range of the LAs (**each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55**);

- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (**column 4, lines 55-59**);
- receiving input/output (IO) requests from host processors directed to specified LAs (**column 2, lines 11-15; column 4, lines 64-67**); and directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).
- wherein a further interim fast-access-time cache is adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails (**paragraph 0044**).

With respect to claim 29, Henry et al. teach the method according to claim 28, wherein the mass storage devices comprise one or more disks (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

With respect to claim 30, Henry et al. teach a system for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) (**Fig. 2; column 4, lines 10-23, lines 55-67**), comprising:

- a plurality of interim fast-access-time caches (**cache pools 1 and 2**), configured to operate as controllers substantially independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (**column 5, lines 45-55**); and coupled to receive data from and provide data to

the one or more -mass-storage devices within the respective second range
(column 4, lines 55-59); and

- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned
(column 2, lines 11-15; column 4, lines 64-67).
- a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails
(paragraph 0044).

With respect to claim 31, Henry et al. teach the storage system according to claim 30, wherein the mass storage devices comprise one or more disks **(disks 1-5 of Fig. 2; column 5, lines 9-23).**

With respect to claim 32, Henry et al. teach a method for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) **(Fig. 2; column 4, lines 10-23, lines 55-67)**, comprising:

- assigning to each of a plurality of interim-fast-access-time caches **(cache pools 1 and 2)**, configured to operate as controllers substantially independently of one another, a respective second range of the LAs **(each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55);**

- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (**column 4, lines 55-59**);
- receiving input/output (IO) requests from host processors directed to specified LAs (**column 2, lines 11-15; column 4, lines 64-67**); and
- directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (**column 2, lines 11-15; column 4, lines 64-67**).
- wherein a further interim fast-access-time cache is adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails (**paragraph 0044**).

With respect to claim 33, Henry et al. teach the method according to claim 28, wherein the mass storage devices comprise one or more disks (**disks 1-5 of Fig. 2; column 5, lines 9-23**).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 9-11 and 22-24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hicken et al. (US 2004/0153727) in view of Karger et al. ("Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web," by in the Proceedings of the 29th ACM Symposium on Theory of Computing, Pages 654-663).

With respect to claim 9, Hicken et al. disclose the method according to claim 1. Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (**page 5, section 4, "Consistent Hashing"**).

Hicken et al. and Karger et al. are analogous art because they are from the same field of endeavor, namely data caching.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the minimum redistribution in the form of consistent hashing of Karger et al. with the data caching redundancy system of Hicken et al. The motivation for doing so would have been because to prevent requiring a central server to distribute a completely updated hash table to all the machines every time a new machine is added to the network (**page 2, column 2, paragraph 2 beginning with "Our second..."**).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Karger et al. with Hicken et al. for the benefit of a data caching system with consistent hashing to obtain the invention as specified in claim 9.

With respect to claim 10, Hicken et al. disclose the method according to claim 9. Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (**page 5, section 4, “Consistent Hashing”**).

With respect to claim 11, Hicken et al. disclose the method according to claim 9. Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (**page 2, column 1, paragraph 5, “Our first tool, *random cache trees*...”**).

With respect to claim 22, Hicken et al. disclose the storage system according to claim 14. Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (**page 5, section 4, “Consistent Hashing”**).

With respect to claim 23, Hicken et al. disclose the storage system according to claim 22. Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (**page 5, section 4, “Consistent Hashing”**).

With respect to claim 24, Hicken et al. disclose the storage system according to claim 22. Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (**page 2, column 1, paragraph 5, “Our first tool, *random cache trees*...”**).

Response to Arguments

Applicant's arguments filed 1 October 2007 have been carefully and fully considered but they are not persuasive.

With respect to Applicant's argument located within the second paragraph of the third page of the remarks (numbered as page 15) which recites:

"However, as is apparent from figure 3 of Hicken, elements 338 and 339 do not operate substantially independently of each other. Rather, if CPU 336 fails, then both primary cache 338 and secondary cache 339 also fail. Therefore, the elements which the Office Action relies on as disclosing the caches of the present invention do not operate substantially independently."

The Examiner respectfully disagrees. Applicant's arguments fail to comply with 37 CFR 1.111(b) because they amount to a general allegation that the claims define a patentable invention without specifically pointing out how the language of the claims patentably distinguishes them from the references. Insofar as the Examiner is having difficulty locating support for Applicant's argument that if CPU 336 fails, then both primary cache 338 and secondary cache 339 also fail. Assuming *arguendo* that Applicant's argument was not just a mere allegation of patentability, the Examiner wishes to draw attention to paragraph 0012 of Hicken that states failure of just the primary cache in a controller. No mention is made that the secondary cache of the controller is rendered inoperable due to such failure. Thus anticipating that which is instantly claimed of the caches as operating "substantially independently."

With respect to Applicant's argument located within the second paragraph of the fourth page of the remarks (numbered as page 16) which recites:

"Furthermore, the Office Action asserts that element 338 discloses the one or more third cache as claimed. However, there is no disclosure relating to reconfiguring element 338 to perform the operations of retrieving data from and storing data at logical addresses which were previously serviced by element 333, which the Office Action asserts discloses the second cache. Nor is there any disclosure that element 338 in Hicken accepts these additional functions while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs, as claimed. Hicken merely discloses

using a backup secondary cache memory for loading data in storage when a primary cache memory fails. Therefore, for at least all of the above reasons, claim 1 is allowable."

The Examiner respectfully disagrees. As noted in Hicken, paragraph 0044, the remaining failed over storage controller searches for an available (lone) storage controller with which to pair-off. The remaining failed over storage controller then pairs with a lone storage controller, creating a pair. The lone storage controller reconfigures its secondary cache that is presently functioning as a redundant cache to its primary cache to act as a 'third' cache as instantly claimed for the failed over storage controller, thus anticipating the instant limitations of a third cache. Furthermore, as taught throughout Hicken, for example in paragraph 0036, the flushing of the caches that is done can be done at different times. That is, after reconfiguring, and before flushing, the cache memory continues to perform at least one of the operations, the storing operation, as instantly claimed.

With respect to Applicant's argument located within the second paragraph of the sixth page of the remarks (numbered as page 18) which recites:

"Hicken does not identically disclose or suggest that the first, second and third caches are implemented in three separate physical units and adapted to function as controllers substantially independently of each other.."

The Examiner respectfully disagrees. First, the Examiner wishes to draw attention to the language, "adapted to" and "configured to" as presently found in the claims. As shown in the MPEP, 2106,

"Language that suggests or makes optional but does not require steps to be performed or does not limit a claim to a particular structure does not limit the scope of a claim or claim limitation. The following are examples of language that may raise a question as to the limiting effect of the language in a claim:

- (A) *statements of intended use or field of use,*
- (B) *"adapted to" or "adapted for" clauses,*
- (C) *"wherein" clauses, or*
- (D) *"whereby" clauses."*

Thus, the language adapted to and configured to does not further limit the claim in that such language represents intended use. Additionally, the Examiner refers Applicant's to the rejection *supra* and the comments made *supra*.

With respect to Applicant's argument located within the last paragraph of the sixth page of the remarks (numbered as page 18) which recites:

"Henry fails to disclose or suggest the feature of a plurality of interim fast-access-time caches, configured to operate as controllers substantially independently of one another, as recited in claim 26."

The Examiner respectfully disagrees. Applicant's arguments, again, fail to comply with 37 CFR 1.111(b) because they amount to a general allegation that the claims define a patentable invention without specifically pointing out how the language of the claims patentably distinguishes them from the references.

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

Art Unit: 2187


the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Matthew Bradley whose telephone number is (571) 272-8575. The examiner can normally be reached on 6:30-3:00 M-F.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald A. Sparks can be reached on (571) 272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

DAS/mb



DONALD SPARKS
SUPERVISORY PATENT EXAMINER